MANAGEMENT OF COMMUNICATION NETWORKS

Field of the Invention

The present invention is concerned with communication networks and more particularly with network management systems that manage various aspects of the communication network. More particularly, the invention concerns a technique for the more efficient management of a communication network that, for various reasons, has been divided or partitioned into subnetworks.

Background to the Invention

A conventional communications network, for example a broadband communications network, comprises a plurality of physical resources in the form of network elements, eg switches, cross-connects, regenerators, repeaters, transmission links such as fibre optic links or coaxial cable links, operating under control of a plurality of logical resources, eg transport protocols, and local controls associated with individual physical resources. An example of a generic representation of a communications network is illustrated schematically in Figure 1, in which the physical resources of a core network are located at a plurality of nodes 100 and links 101 distributed over a geographical area.

For a network operator to maintain control of a communications network for its operation, administration and maintenance, a management system is maintained which stores information describing the physical and logical resources within the network. One or more management systems may reside at a centralized location, eg a network controller 102, or different management systems may be situated at a plurality of network controllers at different locations. The management system stores data describing each individual network element in a communications network and has one or more management applications which use the data to manage various aspects of the network, eg operation, administration, and maintenance of the network.

A conventional communications network may comprise of the order of thousands of individual network elements, eg switches, cross-connects, regenerators, each of which contains of the order of tens to hundreds of cards, having processors, line terminations, buffers, registers, switch fabrics, etc. each card containing of the order of hundreds of individual components. In general, a conventional communications network may comprise a multitude of different legacy equipment types *of different*

proprietary manufacture, each of which has its own particular internal configuration and offers its own specific capabilities.

International Telecommunications Union (ITU-T) recommendation G.805 of November 1995, (available from International Telecommunication Union, General Secretariat, Sales Service, Place de Nation, CH 1211, Geneva 20, Switzerland), sets out a functional architecture for telecommunications transport networks in a technology independent manner. A generic functional architecture is set out as a basis for a harmonized set of functional architecture recommendations for broadband transport network including asynchronous transfer mode (ATM), synchronous digital hierarchy (SDH) and plesiochronous digital hierarchy (PDH), as well as a corresponding set of recommendations for management, performance analysis and equipment specification for such transport networks.

In general, in known transport networks circuit switched communications are made on an end-to-end basis over a plurality of network entities. In this specification, by circuit switched, it is meant that the network reserves part of its resources for the purpose of supporting an end-to-end communication, for the duration of that communication, whether those resources are used or not.

Referring to Figure 2, there is illustrated a simple example of a trail of a circuit switched communication over part of a communications transport network. A transport network is defined in recommendation G.805 as "the functional resources of the network which conveys user information between locations". In recommendation G.805, a trail is defined as "a transport entity which consists of an associated pair of unidirectional trails capable of simultaneously transferring information in opposite directions between their respective inputs and outputs". A unidirectional trail is defined as a "transport entity" responsible for the transfer of information from the input of a trail termination source to the output of a trail termination sink.

The integrity of the information transfer is monitored. It is formed by combining trail termination functions and a network connection. A transport entity is defined as "an architectural component which transfers information between its inputs and outputs within a layer network". A layer network is defined as "a topological component that includes both transport entities and transport processing functions that describe the generation, transport and termination of a particular characteristic infor-

mation". A connection is defined as "a transport entity which consists of an associated pair of uni-directional connections capable of simultaneously transferring information in opposite directions between their respective inputs and outputs". A unidirectional connection is defined as "a transport entity which transfers information transparently from input to output".

In Figure 2, there is illustrated schematically a plurality of transport entities 200, 201, 202, 203, 204 in a communications network comprising network elements eg switches, cross-connects, links, supporting an end to end trail between first and second trail termination points (TTPs) 205, 206. The trail is carried over a plurality of connections which connect the transport entities to each other. Connections between transport entities terminate at a plurality of connection termination points (CTP) within the transport entities.

The generalized trail as illustrated in Figure 2, incorporates different trails in different transport protocols. For example, virtual paths and virtual circuits in asynchronous transfer mode (ATM) constitute trails within the meaning of ITU-T Recommendation G.805. ATM cells may be carried within a virtual path within SDH frames over an SDH trail.

Within a layered network, protocol trails occur within layers. Trails can occur at a plurality of different layers. However, each trail is always contained within its own layer. In a large network, comprising tens to hundreds of network elements, management of end-to-end trails poses a highly complex problem and poses difficulties in the practical implementation of setting up and tearing down of trails. The concept of trail management is mentioned in recommendation G.805 in which a trail management process is defined as "configuration of network resources during network operation for the purposes of allocation, reallocation and routing of trails to provide transport to client networks."

Conventionally, for creating a trail across a network it is known for several network operators, at several network controllers controlling different sections of the network, to each set up one or more connections within sections of the network which they control. To achieve a trail over a large number of transport entities, a network operator wishing to set up a trail may need to contact, by means of a telephone call or a fax, other network operators having control of other parts of the network across

which a trail may pass, and co-ordinate the setting up of a trail by verbal or fax communication with the other human network operators.

In conventional prior art network management systems, it is known to keep a master database which always overwrites whatever connections exist in the real network under management. Thus, if a network operator makes changes to connections or trails in a network by configuring an individual network element directly, the conventional network management system database will attempt to overwrite any changes made at the network element level, regardless of whether the network operator intended those changes to the network element or not. Further, the known network management systems do not provide an ability to draw configuration and connectivity information from the real network and do not compare such information with the information kept in the master database.

Prior art network management systems either represent network configurations which a network operator plans at a network controller, and implements those configurations irrespective of existing configurations of a network, or provide a network operator with data describing actual network configurations, without taking into account or making provision for a network operator's planned or intended present and future configurations of the network.

In the following discussion, a preferred implementation of the invention is described with reference to synchronous digital hierarchy (SDH) systems. However, it will be understood that the scope of the invention is not restricted to SDH systems but extends over any network of physical and logical resources in the telecommunications or computer networks domains having a management information system,

Networks operating asynchronous transfer mode (ATM), synchronous optical network (SONET), integrated service digital network (ISDN) and SDH are specific examples of such networks. However, the invention is not restricted to networks operating these specific protocols.

ITU-T recommendation G.803 deals with the architecture of SDH transport networks and defines an SDH based transport network layered model as illustrated in Figure 3. The G.803 model uses a functional approach to the description of architectures based on the concept of a number of SDH functional layers and the concept of partitioning within a layer for defining administrative domains and boundaries.

Physically, a conventional SDH network is constructed from a plurality of physical resources, for example network elements such as exchanges, multiplexers, regenerators, and cross connects. The network elements are connected together and provide a transmission media layer, including a section layer comprising multiplex section layer 300, a regenerator section layer 301 and a physical media layer 302. Circuit switched traffic is routed over the physical resources in a circuit layer 303 which is carried by the SDH transport layers.

The SDH multiplexing structure is illustrated schematically in Figure 4, which also shows synchronous optical network (SONET) multiplexing options and European Telecommunications Standards Institute (ETSI) multiplexing options. The SDH transport layers comprise, in addition to the physical media layer and section layer, a plurality of higher order path layers, for example carried by virtual containers VC-3, VC-4, and a plurality of lower order path layers, for example carried by virtual containers VC-2, VC-3, VC-11, and VC-12.

Data is carried between network elements that are geographically separated by large distances at relatively high data rates, eg 155 Mbits/s. Circuit switched connections, referred to as a circuit layer 303 in recommendation G.803, are transported across the SDH network by encapsulating bit streams comprising the circuit switched connections into different virtual containers (VCs) which are multiplexed together for transmission at higher order bit rates.

Within the physical resources, circuit switched traffic follows paths and trails at various multiplex levels. Connections are terminated at connection termination points (CTPs), and trails are terminated at trail termination points (TTPs) within physical resources. For example, within a communications network, there may be a restricted number of network elements that are capable of processing voice data.

Operations on voice data at a voice level may be performed within those particular network elements. However, to transport traffic data between those network elements, there must be further transmission, such as provided by the SDH virtual container system. Thus, where a voice connection is to be made between geographically disparate network elements A and B, the connection may be routed via intermediate network elements D, E, F, G etc which may be in the VC-12 layer. However,

the VC-12 layer itself, to connect between intermediate network elements E, F, may need to be multiplexed into a higher bitrate layer, eg the VC-4 layer.

Referring to Figure 5, there is illustrated schematically a section of an SDH communications network comprising a plurality of network elements 500—505 operating under control of an element controller 506 and managed by a network controller, referred to herein as network resource manager 507.

The element controller communicates with the plurality of network elements via an operations administration and control channel 509, eg using a conventional network management protocol, for example the known common management information service element (CMISE) protocol. The element controller communicates with the network resource manager 507 via a conventional protocol, for example the transmission control protocol/internet protocol (TCP/IP) over a transmission link 508. The network resource manager 507 implements control of the network by implementing operations, administration and management operations of the network elements, through one or a plurality of element controllers 506.

Referring to Figure 6, there is illustrated schematically the construction of a typical network element 600, element controller 506 and network resource manager 507. Network element 600, for example a multiplexer or cross connect, comprises a casing or cabinet having one or a plurality of shelves, each shelf containing a plurality of cards 601. The cards contain processors, switch fabrics, line terminations etc depending upon the type of network element, and are connected to each other via a data bus. In the case of an SDH multiplexer, each card may support a number of physical ports. Each port supports a plurality of connections. The network element is provided with a local control system 602 comprising a data processing capability configured to send and receive messages over the CMISE OAM channel 509.

The element controller comprises a workstation 603, for example a Hewlett Packard 9000 series workstation comprising a processor 604, a data storage device 605, a bus 606 linking the processor and data storage device, a graphical user interface 607, and a communications port 608 for communicating with the network element and the network resource manager. Typically, the element controller operates according to a UNIX operating system 609.

The network resource manager 507 similarly may comprise a work station 610, eg Hewlett Packard 9000 series having processor 611, memory 612, bus 613, graphical user interface 614 and communications ports 615 components, operating in accordance with a UNIX operating system 616. The network resource manager and the element controller are configured to communicate with each other using for example TCP/IP link 508.

The network resource manager comprises a managed object base (MOB) 617 containing data describing characteristics and configurations of the network elements under its management. Within the net work resource manager, each network element is represented as a managed object, in accordance with the telecommunications network management network (TMN) architecture of ITU-T recommendation M.3010.

In managed object base 617 physical resources of the network, comprising the transport entities supporting the trails, eg central office switches, multiplexers, regenerators, cross-connects etc are represented as managed objects according to ITU-T recommendation M.3010 (Principals for a Telecommunications Management Network) in known manner. Additionally, individual capabilities and functionalities of those physical resources, for example trail termination points, connection termination points and adaptations within individual physical or logical ports of the physical resources, and the connection limitations and connectivity capabilities of those physical resources are represented within managed object base 617 according to an object representation scheme as disclosed in co-pending US Patent Application Serial No 09/010,387 (corresponding to EP 98306103.7) entitled "Capability Modeling Using Templates in Network Management System".

The network resource manager 507 comprises a trail manager application 620 for managing trails across the network. Management operations controlled by trail manager application 620 are implemented at each of a plurality of element controllers 506 by respective trail management operation controller server 619. Trail manager application 620 provides a network operator with means for managing trails across a network. In order to enable an operator to manage trails, trail manager application 620 is provided with functionality for:

planning trails across the network;

learning about actual existing trails within the network;

storing data describing existing trails within the network provisioned from planned trails; and

storing data describing whether a planned or provisioned trail is intended within the network.

Within a communications network, although a network operator may create and manage trails using trail management application 620, actual trails may exist within the network which are different to those which the network operator has intended. Trail management application 620 is provided with a graphical user interface (GUI) 614 which enables the network operator to view both the actual trails within the network and the network operator's planned and/or intended trails within the network. For each trail under management of the trail management application 620, there is maintained data representing a status of the trail. The means for representing the status of each trail comprises a state machine that is part of the trail manager application 620, providing data to the trail manager application.

The state machine comprises data processing capability and data storage capability (a database) for maintaining and processing data describing one or more states of each trail under management. In the specific implementation herein, the state machine is physically implemented by configuration of the processing and data storage capabilities of the conventional network management system, for example one or more Hewlett Packard 9000 Series Workstations configured as the element controller, and network resource manager as illustrated in Figure 6.

Such configurations are implemented by arrangement and allocation of a data storage device and by provision of a set of algorithms to perform data processing operations on data stored on the database. Such arrangements and algorithms may be implemented in a conventional programming language, such as the known C/C++ language as will be appreciated by those skilled in the art. Specific programming options and variations of implementations are numerous and will be readily apparent to the skilled person.

The trail manager 620 obtains data describing specific trail termination points within individual network elements, from managed object base 617, as described in the aforementioned co-pending patent application, and is thereby provided with information concerning available capacity and connection capabilities for supporting

trails and connections. The trail manager application 620 obtains data describing the capabilities, including connectivities and restrictions on connectivities of each of the network elements by referencing a set of data templates stored in the managed object base. The templates include templates describing physical or logical ports of a network element, together with connection templates describing possible connectivities of termination points within each physical or logical port of a network element on a layer by layer basis.

In general, each port supporting a trail is represented by a column of layers, as illustrated in Figure 8. Depending upon the protocol layers supported by the ports, the height of the column may differ from port to port. Figure 8 illustrates schematically a data representation of part of the VC-12 trail over network elements 700, 701 as stored in the managed object base 617. For each network element, a physical or logical port supporting the trail is represented as an assembly of termination point data templates 900, represented by symbols as illustrated in Figure 9. Symbol 901 represents a trail termination point, symbol 902 represents an adaptation between a same layer of the trail termination point and a client layer, symbol 903 represents connectivity to a client layer, and symbol 904 represents connectivity to other termination points in the same layer.

In Figure 8, a trail, eg a VC-12 trail, enters first network element 700 at VC-12 termination point 703 through VC-12 adaptation 800 at a first port 801 of first network element 700. Transport between first and second network elements over link 705 is effected over SDH physical media section 802 to an entry port 803 of second network element 701. Conversion of the physical media section through the SDH protocol layer is represented by a set of data templates representing the physical media section layer 802, optical section layer 805, regenerator section layer 806, STM-N layer 807 and HP-VC4 layer 808, each represented by a separate data template as illustrated in Figure 9. Internal connections between input and output ports 803, 804 within same network element 701 is made via a VC-12 connection 712.

Referring to Figure 10 herein, a trail 1000 between trail termination source point 703 and trail termination sink point 704 may be set up by a network operator at network resource manager 506, similarly as described in Figure 7 herein. The trail manager 620 has a record of the actual trail in the network as described with reference

to Figure 7 herein, from data read from managed object base 617, in accordance with the data template representations described with reference to Figures 8 and 9 herein. However, in the network, the actual trail may become altered from that created or intended by the network operator, for various reasons. For example, maintenance personnel may be able to take local control of network elements in order to reconfigure connections directly at the network element level, overriding the network resource manager 506 and element controller 507. Thus, in this example in practice an actual trail may be reconfigured, due to local alterations made at second network element 701 so that the VC-12 trail is re-routed to a fourth network element 1001 as shown in Figure 10. Thus, a new actual trail 1002 exists in the network between second trail termination source point 1003, through fourth network element 1001, second network element 701, and third network element 702 o end at trail termination sink point 704. Therefore, whilst a network operator at network resource manager 507 intends a first trail between first and third network elements as shown in Figure 7, due to external circumstances beyond the network operator's control, eg due to local reconfiguration of second network element 701, an actual trail between fourth and third network elements may be created as illustrated in Figure 10 herein, which is different to the intended first trail, and overwrites it.

In many cases, the actual trails within the network are the same as trails intended by the network operator. However, discrepancies between intended and actual trails do occur. To provide comprehensive trail management throughout the network, the state machine keeps a record of:

planned trails, eg as input by a network operator at GUI 714 of network resource manager 507; and

actual trails within the network, eg created at network resource manager 507 and provisioned in the network, or as a result of events occurring within the network independently of network resource manager 507 and element controller 506.

Planned and actual trails may either be intended or unintended. Usually, the intention of a network operator is that all trails planned at the network resource manager 507 become executed as actual provisioned trails in the network. However, trails which were not planned at the network resource manager may or may not be intended.

In the case of the example of Figure 10, trail manager application 620 records the actual trail 1002 between fourth and third network elements, the intended trail 1000 between first and third network elements, and the fact that the actual trail 1002 between fourth and third network elements may also be intended (since it is a valid trail). Additionally, the state machine may record data representing that each trail identified in Figure 10 is a valid trail, and that they are in conflict: that is to say both trails cannot exist at the same time in the network, because they are mutually exclusive in terms of their demands on the network elements, as well as recording which of the trails was originally planned, and which of the trails has been learnt from interrogation of the network, and may indicate that the trail manager application 620 cannot resolve the discrepancy between the two trails.

The state machine maintains one or more state models for each trail under management of the trail manager 620. The trails may be either actual trails existing within the network, or trails intended to be created or having been created by the network operator. A state model comprises a data record of a trail which records a state in which the trail currently resides, ie a condition of the trail. The data is held in a database containing a list of trails within a network, together with data describing the status and characteristics of the trail according to a state model.

For each trail there is maintained data describing the trail in managed object base 620 in the form of one or a set of trail objects. The state machine performs automatic operations on the trail objects, depending upon which one of a plurality of possible states they reside. The state machine carries out automatic processes, such as those relating to provisioning, supporting, and creating database trails. Further, a network operator may activate operations on the trail objects, eg by typing in commands at graphical user interface 614 (Fig 6).

Network management presents certain particular problems, especially with the escalation in size and proliferation of networks. In order to make network management more manageable, networks are often partitioned into a number of subnetworks for a number of different, and often disparate, reasons, such as those considered below:

First, in multi-vendor networks, where the network management system does not manage all of the equipment in the network;

Second, exceptionally large networks cannot be managed effectively by a single network manager and partitioning is essential to deal with the network in small, cohesive chunks;

Third, interworking with a subnetwork control plane management system, such as ASTN subnetworks, carries a requirement that the network management system is able to partition into controlled subnetworks;

Fourth, there are occasionally requirements for a network to be partitioned in such a way that some control can be offered over a portion of the network without affecting the rest of the network (Customer Network Management); and

Fifth, the operations of a network or regional implication indicate that the network must be managed in partitioned subnetworks, from a number of perspectives (multi-jurisdictional management).

Previous attempts at managing communication networks have involved setting up "dummy" or "pseudo" network elements to simulate other parts of the network. The apparent simplicity of this prior "solution" actually breaks down because of the difficulties in managing such a partitioned network since it introduces a configuration and assignment task to manage the dummy/pseudo nodes/network elements. Moreover, this known approach does not solve all of the problems associated with network management. In particular, it does not enable partial or incomplete trails.

If that were not enough, the resultant network becomes unwieldy, requires far too much administration and overhead, and adds unnecessary complexity. Ideally, a solution should be transparent as far as the customer is concerned, whereas previous solutions involve the customer more than is welcome and/or desirable.

The invention therefore provides a solution that does not suffer from the disadvantages of prior approaches to solving the problems of partitioned network management.

Summary of the Invention

In its broadest sense, the invention is based on the recognition that effective management of partitioned networks can be achieved by introducing an "off-network" pointer to enable a traffic carrying capability to be established to locations off network.

In a first aspect, the invention provides a method of managing a communication network comprising a plurality of ports, modelled according to a layer protocol, and a network management system, the communication network being partitioned into a plurality of subnetworks, the method comprising generating, in respect of a said subnetwork, an off-network pointer exiting the subnetwork at one of said ports, whereby to establish a traffic carrying capability externally to the subnetwork.

The pointer is preferably first generated at the physical layer and functionality at higher layers are generated in response so as to generate a high order transport connection to carry said traffic.

In a second aspect, the invention provides a method of managing a communication network comprising a plurality of ports, modelled according to a layer protocol, and a network management system, the communication network being partitioned into a plurality of subnetworks, the method comprising determining those ports that represent connection termination points in the subnetworks, whereby to generate trails interconnecting said connection termination points in different subnetworks.

In each case, the method may further comprise identifying incomplete trails within a partition to identify stranded bandwidth and to create trails piecemeal.

The invention also provides a communication network, a user interface and a network management system employing the same techniques as in the above methods. The invention is preferably carried out by software and includes communication signals transported over the network. Finally, the invention includes a carrier carrying the software.

The preferred features as described hereinabove or as described by any of the dependent claims filed herewith may be combined, as appropriate, as would be apparent to the person skilled in the art, and may be combined with any of the aspects of the invention as described hereinabove or as described by any of the independent claims filed herewith.

Reference is directed to co-pending US Patent Application entitled “Communication Networks” (Nortel Networks Limited’s file reference 14568IDUS01U), a copy of which is annexed to the present application for incorporation herewith, for further details of the manner in which connectivity information at one layer of a

communication protocol can be used to derive connectivity information at another layer or layers.

Brief Description of the Drawings

The invention will now be described with reference to the drawings, in which:

Figure 1 is a generic representation of a communications network in which the present invention may operate;

Figure 2 is a simple example of a trail of a circuit switched communication over part of a communications transport network;

Figure 3 illustrates an SDH based transport network layered model;

Figure 4 illustrates an SDH multiplexing structure;

Figure 5 illustrates schematically a section of an SDH communications network;

Figure 6 illustrates schematically the construction of a typical network element;

Figure 7 illustrates schematically a generic representation of a trail;

Figure 8 illustrates schematically a plurality of ports comprising network elements of a network;

Figure 9 illustrates schematically a data representation of a port;

Figure 10 illustrates schematically a planned trail and an existing trail within a network;

Figure 11 represents schematically a network incorporating the invention;

Figure 12 represents the building up of connectivity at different protocol layers in connection with an off-network trail;

Figure 13 illustrates a partitioned region in a network with off-network pointers;

Figure 14 represents the protocol layers applied to an off-network pointer;

Figure 15 illustrates the relationship between Trail/TTP and Link/LTP in a network;

Figure 16 illustrates the invention as applied to a CTP-terminated region; and

Figure 17 illustrates the Trail/CTP relationship.

Detailed Description of the Illustrated Embodiments

The above introduction presented the environment in which implementations of the present invention are intended to be placed. It is now necessary to consider particular situations within that environment in order to appreciate the significance of the invention and its contribution towards improved management of a communications network. Three scenarios are contemplated, namely: off-network, CTP terminated and other.

Off-network

Figure 11 represents schematically a communications network in which an OSS (Operation Support System) 1100 is in overall command of a plurality of network management systems (NMSs) 1101, 1102 ...110n. Each NMS is associated with one or more Management Domains (MDs), such as 1103 to 1106. One such MD 1103 is shown associated with a first region R1 indicated at 1107. Other MDs are associated with other regions R2, R3 and so on, as indicated at 1113, 1114 etc.

Each region represents a geographical area that is modelled on the network. However, modelling is not possible where one (or more) of the regions is not owned by the network provider. It also follows that, for every region not owned by the network provider, there can be no management of that region. Within each region there will be numerous NEs, such as those indicated generically at 1110, 1111, 1112, 1113, 1114 and 1115. Connections within a region are made between the "internal" NEs. Connections between regions can only occur between NEs built to the edge, in modelling terms, such as 1110, 1111 and 1112. In terms of network modelling, there will be pointers leading out of the regions R1, R2 from the NEs 1111, 1112 to common interconnection systems, such as via GPS references or via Graphical Information Systems.

Assume, for the sake of non-limiting example, that MD 1103 and "its" region R1 are not owned by the network provider operating MD 1104 and "its" region R2. As far as region R2 is concerned, a trail leaving R2 is under management of the NMS for that region up to the edge but not beyond. A termination point on NE 1110 and an NE 1111 are assumed to be at edges of region R1. Similarly, NE 1112 is assumed to be at the edge of region R2 whereas NE 1113 is somewhere within the region. Again, NE 1114 is within region R3. NEs 1112 and 1113 may be part of a photonic sub-network whereas NE 1114 may be within a SDH/SONET sub-network. Other sub-

networks or combinations of sub-networks, and data formats, such as physical, PDH/Async, FICON/ESCON/1Gb ethernet or ATM/FR/IP/MPLS etc, may all be handled by the invention.

In a network region or domain such as R1, under the control of one operator and carrying traffic across it between points 1110 and 1111 and between one or more different operators outside the control of the first operator, there may be an optical fibre traversing the region R1 and going off-domain to the other operator regions, such as R2. Although this discussion is centred on the physical layer, it is to be noted that the invention is equally applicable to the other, logical layers in the hierarchy.

In Figure 13, the NEs corresponding to 1110 and 1111 are indicated as Y and X respectively. In terms of modelling the network, at the physical layer (in this example) the physical trail external to NE X only has one end since the network operator only has information about the trail within its own network. The trail in question is illustrated as extending from the trail termination point (TTP) of the network element in question and is designated as going off-domain, where there is no information about connectivity etc. However, some attribute information can be assigned to the unknown end so as to indicate where the trail is connected to.

The assumption is made that the unknown other end will be compatible with the TTP at the known end within the region under the operator's control. It is a reasonable assumption since, if the other end were not compatible, there would be little prospect of communication between the regions via these elements. It is then assumed that there is an equivalent adapter at the other end under control of the other operator. This assumption is then extrapolated up through the OS, Regeneration and Multiplex layers, as illustrated in Figure 12. Eventually, a higher order layer, such as VC4 is reached where there may be some flexibility.

The node X at this point may be able to cross-connect so as to perform a useful function as regards cross-region traffic of VC4 characteristics. For example, a traffic carrying capability could be built across the network, as indicated by the broken arrow in Figure 13. To summarise, from an assumption as to the physical layer connectivity, the functionality at the logical layers can be deduced and it becomes possible to build a service from node Y, for example, within the region, to node X which is off-network, thereby effectively terminating the trail off-network. The net-

work can thus be built to well-defined points. Without this capability, it would not be possible to differentiate between other nodes within the region so as to determine those that may be used legitimately to transport traffic across the region between certain ports. The connectivity of some ports may not yet be completely defined. The invention therefore enables the building of services that go off-network or off-domain without disrupting the basic G.805 model.

This is achieved in the following manner. In the same way as a trail has a TTP at one end but not at the other because it is undefined, a link is instantiated in a similar frame having an LTP at one end and nothing at the other end. The link offers a service on which to build out into the network. This can be extended further. Rather than simply providing transport through the region, the invention enables a service to be built within the region under the operator's control into another region under a different operator's control. For example, a VC-12 service can be built onto the VC4 service indicated above. From information about off-network connectivity at the physical or logical layers, the invention enables logical functionality to be extrapolated up to the point where a "logical pipe" extends across the region to another operator's region, enabling a service to be built across to the other region. The point is ultimately reached where services are offered that can in turn offer services.

In this way, transport to other regions can be provided from not just the edge nodes/Network Elements but from nodes interior to a region/domain. This is made possible because the logical implications of off-network pointers on client layers/protocols are determined.

Figure 12 illustrates the manner in which the traffic carrying capability is built up from the physical layer, (as shown at the left hand side of the Figure) and is extrapolated into the other, logical layers, eventually reaching the upper container layers with VC4 or VC12 capability, for example. In this way, the elementary connectivity information available at the edge of the region is "converted" into logical functionality information that enables the trail to be built back from the edge and firmly into the inner areas of the region so as to provide the traffic carrying capability into other region(s). There may be a number of units indicated "n.." in the Figure.

A trail or a link normally has two ends. In the present modelling scenario, one end would be a real TTP, the other end is essentially unknown but would have a ter-

mination point of some kind. TTPs terminate the signal (they extract the data header and bring out the contents); LTPs represent the termination point offered to the client layer and are indicative of the potential to offer service to that client layer. LTPs represent the TTPs but in the client layer, with understanding of the client layer's characteristics, eg bandwidth. CTPs represent the utilisation of a particular channel. CTPs inspect the data headers but do nothing with the header or the content. They merely pass them on to the next CTP in the path, ultimately ending in a TTP.

The invention is not restricted to one end having a TTP. In fact, both ends may well be off-network in the case of a carrier's carrier network where transport services are offered unterminated through a region between other regions. In this scenario it is not known what client layers/protocols are supported as there are no termination or adaptation functions known in the managed region.

All three of these termination point types/classes (TTP, LTP, CTP) are derivations of a base termination point type/class and inherit (in C^{++} terms) properties from the TP such as annotations and references. See Figure 14. The termination point TP can be annotated, for example with an indication as to the piece of equipment it represents/is going to, where the fibre needs to be connected. The same concept applies to links. These normally connect two LTPs together but, in this case, there can be an LTP at one end and a TP at the other. The trail to link relationship is analogous to the TTP to LTP relationship, where the link represents the trail within the layer. In this manner, the trail in the present scenario has offered a link to the VC12 layer.

The invention does not impede protection that is normally provided in communication networks. It enables protected clients to be built on off-network trails.

CTP Terminated

In this scenario, as depicted in Figure 16, one large network is envisaged, under the control of one operator and partitioned into a number of domains of control for a variety of reasons, for example from an organisational perspective, administratively, for reasons of network technology characteristics such as ASTN and so on. The objective is to build an end-to-end service under theses circumstances. The individual partitions may represent different geographical and/or administrative areas. However, it is important in this scenario that off-domain or off-network pointers are not used since the overall network, although partitioned, remains a single network *under the*

control of one operator. For this reason, instead of building to an off-network pointer, the trails are built to CTPs. This enables the true topology of the network to remain intact.

Consider the situation that the operator of the backbone of the network has been requested to provide a service through the backbone to provide transport between a number of different points. In this case, the trail would reference a CTP, rather than a TP, at Y in Figure 16, at an edge of the backbone. There may be a CTP at each end, as shown at Y, X in Figure 16.

The operators in the partitioned areas can "see" up to and including the demarcation points, such as X, Y, between partitions. When the physical layer is built between Y, say, in the backbone and C in an adjacent partition, and between X in the backbone and F in another partition, it is possible to support connectivity between regions. To achieve this it is necessary to be able to define the demarcation points, ie the CTPs that are members of a particular domain and are edge points. It is undesirable to allow trails to be built to any points within partitions as this leads to a lack of control and hinders interpretation of the intended trail as it would no longer be possible to identify when trails were validly complete. The invention enables trails to be built that start on a TTP or a CTP on the edge of a partition and build over a number of other NEs and then terminate on a TTP or a CTP at another edge. Trails are built from say A through B and C to Y, from F to X and from G to U, for the sake of example. The points A, F and G may go off to geographically distinct areas of a single network. The CTP terminated aspect of the invention enables the capability to build end-to-end trails in parts according to the partitioning needs of the operator. For example, connectivity can be completed in the backbone by building a trail/circuit fragment from Y to X via U to provide transport between regions through the backbone. The invention also enables access to be built between different domains in the network for administrative or technological reasons such as access to an ASTN capable domain.

Incomplete trails

When connectivity is discovered in a network, it is important to identify network services that are meaningful and are meant to exist or not to exist. This is done by defining the partition points, ie whether they are off-network or CTP terminated.

If this identification is not possible, the conclusion is made that the trail, such as between points 1 and 2 in Figure 16, is incomplete as terminating points on NEs 1 and 2 are not designated as off-network or valid CTP partitioning points. If so, this probably represents "stranded bandwidth" that is fulfilling no useful function and the cross-connects used to transport the bandwidth are tied up unnecessarily.

This aspect of the invention also caters for the situation where a trail is in the process of being built but is halted for some reason. At that stage, the incomplete trail is committed but will not enter the network since it is incomplete. In this way, trails can be built piecewise.